

dlvhex: A Tool for Semantic-Web Reasoning under the Answer-Set Semantics*

Thomas Eiter, Giovambattista Ianni, Roman Schindlauer, and Hans Tompits

Institut für Informationssysteme 184/3, Technische Universität Wien,
Favoritenstraße 9–11, A-1040 Vienna, Austria
{eiter, ianni, roman, tompits}@kr.tuwien.ac.at

Abstract. We briefly report about the development status of dlvhex, a reasoning engine for HEX-programs, which are nonmonotonic logic programs featuring both higher-order atoms as well as external ones. Higher-order features are widely acknowledged as useful for various tasks and are essential in the context of meta-reasoning. Furthermore, the possibility to exchange knowledge with external sources in a fully declarative framework such as answer-set programming (ASP) is particularly important in view of applications in the Semantic-Web area. Through external atoms, HEX-programs can deal with external knowledge and reasoners of various nature, such as RDF datasets or description-logic bases.

1 Introduction

Nonmonotonic semantics is often requested by Semantic-Web designers in cases where the reasoning capabilities of the *Ontology Layer* of the Semantic Web turn out to be too limiting, since they are based on monotonic logics. The widely acknowledged answer-set semantics of nonmonotonic logic programs [5], which is arguably the most important instance of the *answer-set programming* (ASP) paradigm, is a natural host for giving nonmonotonic semantics to the *Rules*, *Logic*, and *Proof Layers* of the Semantic Web.

However, for important issues such as *meta-reasoning* in the context of the Semantic Web, no adequate answer-set engines have been made available so far. Motivated by this fact and the observation that, furthermore, interoperability with other software is an important issue (not only in this context), in previous work [3], the answer-set semantics has been extended to HEX *programs*, which are *higher-order logic programs* (which accommodate meta-reasoning through *higher-order atoms*) with *external atoms* for software interoperability. Intuitively, a higher-order atom allows to quantify values over predicate names, and to freely exchange predicate symbols with constant symbols, like in the rule $C(X) \leftarrow \text{subClassOf}(D, C), D(X)$. An external atom facilitates the assignment of a truth value of an atom through an external source of computation. For instance, the rule $t(\text{Sub}, \text{Pred}, \text{Obj}) \leftarrow \&RDF[\text{uri}](\text{Sub}, \text{Pred}, \text{Obj})$ computes the predicate t taking values from the predicate $\&RDF$. The latter predicate extracts RDF

* This work was partially supported by the Austrian Science Fund (FWF) under grant P17212-N04, and by the European Commission through the REVERSE IST Network of Excellence (IST-2003-506779).

statements from the set of URIs specified by the extension of the predicate *uri*; this task is delegated to an external computational source (e.g., an external deduction system, an execution library, etc.). External atoms allow for a bidirectional flow of information to and from external sources of computation such as description-logic reasoners. By means of HEX-programs, powerful meta-reasoning becomes available in a decidable setting, e.g., not only for Semantic-Web applications, but also for meta-interpretation techniques in ASP itself, or for defining policy languages.

Other logic-based formalisms, like TRIPLE [10] or F-Logic [8], feature also higher-order predicates for meta-reasoning in Semantic-Web applications. Our formalism is fully declarative and offers the possibility of nondeterministic predicate definitions with higher complexity in a decidable setting. This proved already useful for a range of applications with inherent nondeterminism, such as ontology merging (cf. [11]) or match-making, and thus provides a rich basis for integrating these areas with meta-reasoning.

2 HEX-Programs

2.1 Syntax

HEX-programs are built on mutually disjoint sets \mathcal{C} , \mathcal{X} , and \mathcal{G} of *constant names*, *variable names*, and *external predicate names*, respectively. Unless stated otherwise, elements from \mathcal{X} (resp., \mathcal{C}) are written with first letter in upper case (resp., lower case), and elements from \mathcal{G} are prefixed with “&”. Constant names serve both as individual and predicate names. Importantly, \mathcal{C} may be infinite.

Elements from $\mathcal{C} \cup \mathcal{X}$ are called *terms*. A *higher-order atom* (or *atom*) is a tuple (Y_0, Y_1, \dots, Y_n) , where Y_0, \dots, Y_n are terms and $n \geq 0$ is its *arity*. Intuitively, Y_0 is the predicate name; we thus also use the familiar notation $Y_0(Y_1, \dots, Y_n)$. The atom is *ordinary*, if Y_0 is a constant. For example, $(x, rdf:type, c)$ and $node(X)$ are ordinary atoms, while $D(a, b)$ is a higher-order atom. An *external atom* is of the form

$$\&g[Y_1, \dots, Y_n](X_1, \dots, X_m), \quad (1)$$

where Y_1, \dots, Y_n and X_1, \dots, X_m are two lists of terms (called *input list* and *output list*, respectively), and $\&g$ is an *external predicate name*.

It is possible to specify *molecules* of atoms in F-Logic-like syntax. For instance, $gi[father \rightarrow X, Z \rightarrow iu]$ is a shortcut for the conjunction $father(gi, X), Z(gi, iu)$.

HEX-programs are sets of rules of the form

$$\alpha_1 \vee \dots \vee \alpha_k \leftarrow \beta_1, \dots, \beta_n, \text{not } \beta_{n+1}, \dots, \text{not } \beta_m, \quad (2)$$

where $m, k \geq 0$, $\alpha_1, \dots, \alpha_k$ are higher-order atoms, and β_1, \dots, β_m are either higher-order atoms or external atoms. The operator “not” is *negation as failure* (or *default negation*).

2.2 Semantics

The semantics of HEX-programs is given by generalizing the answer-set semantics [3]. The *Herbrand base* of a program P , denoted HB_P , is the set of all possible ground versions of atoms and external atoms occurring in P obtained by replacing variables with

constants from \mathcal{C} . An *interpretation relative to P* is any subset $I \subseteq HB_P$ containing only atoms.

We say that an interpretation $I \subseteq HB_P$ is a *model* of an atom $a \in HB_P$ iff $a \in I$. Furthermore, I is a model of a ground external atom $a = \&g[y_1, \dots, y_n](x_1, \dots, x_m)$ iff $f_{\&g}(I, y_1, \dots, y_n, x_1, \dots, x_m) = 1$, where $f_{\&g}$ is an $(n+m+1)$ -ary Boolean function associated with $\&g$, called *oracle function*, assigning each element of $HB_P \times \mathcal{C}^{n+m}$ either 0 or 1 (i.e., *false* or *true*, respectively).

This definition of satisfaction, together with a modified notion of a *reduct* as defined by Faber *et al.* [4], enables us to define a conservative extension of the answer-set semantics for HEX-programs. For more details, cf. [3].

Note that the answer-set semantics may yield multiple models (i.e., answer sets) in general. Therefore, for query answering, *brave* and *cautious reasoning* (truth in some resp. all models) is considered in practice, depending on the application.

2.3 Usability of HEX-Programs

An interesting application scenario, where several features of HEX-programs come into play, is *ontology alignment*. Merging knowledge from different sources in the context of the Semantic Web is a crucial task [2] that can be supported by HEX-programs in various ways:

Importing external theories. This can be achieved as in the following manner:

$$\begin{aligned} triple(X, Y, Z) &\leftarrow \&RDF[uri](X, Y, Z), \\ triple(X, Y, Z) &\leftarrow \&RDF[uri2](X, Y, Z), \\ proposition(P) &\leftarrow triple(P, rdf:type, rdf:Statement). \end{aligned}$$

Searching in the space of assertions. In order to choose nondeterministically which propositions have to be included in the merged theory and which not, statements like the following can be used:

$$pick(P) \vee drop(P) \leftarrow proposition(P).$$

Translating and manipulating reified assertions. E.g., it is possible to choose how to put RDF triples (possibly including OWL assertions) in an easier manipulable and readable format, and to make selected propositions true such as in the following way:

$$\begin{aligned} (X, Y, Z) &\leftarrow pick(P), triple(P, rdf:subject, X), triple(P, rdf:predicate, Y), \\ &\quad triple(P, rdf:object, Z), \\ C(X) &\leftarrow (X, rdf:type, C). \end{aligned}$$

Defining ontology semantics. The semantics of the ontology language at hand can be defined in terms of entailment rules and constraints expressed in the language itself or in terms of external knowledge, like in

$$D(X) \leftarrow subClassof(D, C), C(X) \quad \text{and} \quad \leftarrow \&inconsistent[pick],$$

where the external predicate $\&inconsistent$ takes a set of assertions as input and establishes through an external reasoner whether the underlying theory is inconsistent.

Performing default and closed-world reasoning. Assuming that a generic external atom $\&DL[C](X)$ is available for querying the concept C in a given description logics base, the *closed-world assumption* (CWA) can be stated as

$$C'(X) \leftarrow \text{not } \&DL[C](X), \text{concept}(C), \text{cwa}(C, C'),$$

where $\text{concept}(C)$ is a predicate which holds for all concepts and $\text{cwa}(C, C')$ states that C' is the CWA of C , i.e., each individual not explicitly found in C should be in C' .

Inconsistency of the CWA can be checked by pushing back inferred values to the external knowledge base:

$$\begin{aligned} \text{set_false}(C, X) &\leftarrow \text{cwa}(C, C'), C'(X), \\ \text{inconsistent} &\leftarrow \&DL1[\text{set_false}](b), \end{aligned}$$

where $\&DL1[N](X)$ effects a check whether a knowledge base, augmented with all negated facts $\neg c(a)$ such $N(c, a)$ holds, entails the empty concept \perp (entailment of $\perp(b)$, for any constant b , is tantamount to inconsistency).

3 Implementation

The evaluation principle of *dlvhex* is to split the program according to its dependency graph into components and alternately call an answer-set solver (DLV [9]) and the external atom functions for the respective subprograms. The framework takes care of traversing the tree of components in the right order and combining their resulting models. Composing the initial dependency graph from a nonground program is not a trivial task, since higher-order atoms as well as the input list of an external atom have to be considered. To this end, we defined a novel notion of *atom dependency*, which extends the traditional understanding of dependencies within a logic program. This leads to novel types of *stratification* which help splitting a HEX-program and choosing the suitable model generation strategies.

Further methods of increasing the efficiency of computation include a general classification of external atoms regarding their functional properties. For instance, their evaluation functions may be *monotonic* or *linear* (in the sense of a linear function) with respect to a given input. Formalizing such knowledge allows for an intelligent caching algorithm and thus for a reduction of interactions with the external computation source. Latest developments also include a directive to syntactically handle namespaces and an algorithm for traversing the component graph for disjunctive programs, eventually implementing the full HEX-program semantics.

To keep the development and usage of external atoms as flexible as possible, we decided to embed them into *plug-ins*, i.e., libraries that define and provide one or more external atoms. Such plug-ins are implemented as shared libraries, which link dynamically to the main application at runtime. A lean, object-oriented interface reduces the effort of developing custom plug-ins to a minimum.

Currently, *dlvhex* provides the following extension to pure HEX-reasoning: (i) parsing both templates as well as frame syntax by using DLT [7] as a parser; (ii) in addition to strict constraints, accepting *weak constraints* for optimization problems; and (iii) returning the result in XML syntax according to the RuleML specification [1].

The following external atoms are available in dlhex:

The RDF plug-in. The RDF plug-in provides a single external atom, the *&rdf*-atom, which enables the user to import RDF-triples from any RDF knowledge base. It takes a single constant as input, which denotes the RDF-source (a file path or Web address). The *&rdf*-atom interfaces the Raptor RDF library.

The description-logic plug-in. To query description-logic knowledge bases, we developed the description-logic plug-in, which includes four external atoms, allowing for extending a description-logic knowledge base before submitting a query, by means of the atoms' input parameters:

- the *&dlC* atom, which queries a concept (specified by an input parameter of the atom) and retrieves its individuals,
- the *&dlR* atom, which queries an object property and retrieves its individual pairs,
- the *&dlDR* atom, which queries a datatype property and retrieves its pairs, and
- the *&dlConsistent* atom, which tests the (possibly extended) description-logic knowledge base for consistency.

The description-logic plug-in can access OWL ontologies, i.e., description-logic knowledge bases in the language *SHOIN(D)*, utilizing the RACER reasoning engine [6].

The string plug-in. For simple string manipulation routines, we provide the string plug-in. It currently consists of five atoms:

- the *&concat* atom, which lets the user specify two constant strings in the input list and returns their concatenation as a single output value,
- the *&strstr* atom, which tests two strings for substring inclusion,
- the *&split* atom, which splits a string along a given delimiter and retrieves a specific part,
- the *&cmp* atom, which lexicographically compares two strings, and
- the *&sha1sum* atom, which calculates a SHA1 160-bit checksum for a given string.

The policy plug-in. The policy plug-in was created to satisfy the needs of optimization problems that cannot be tackled using conventional methods such as weak or weight constraints in an intuitive way. In the area of policy specification, answer-set programming is used to generate a search space of valid combinations of credentials, which then need to be ranked based on the specific selection of credentials in each solution. For instance, credentials might have various levels of sensitivity regarding their publication in a business transaction, and the overall goal is to find a set of credentials with minimum overall sensitivity. As soon as this overall value is composed in a more complicated way than just the sum of all single sensitivity values, the *&policy*-atom provided by the policy plug-in offers a natural solution. It takes a single predicate as input and returns a numerical value, which is computed according to the predicate's extension and a custom function, implemented by the program designer. Using this value in a single weight constraint facilitates the compact formulation of such an optimization task.

The following code fragment illustrates this technique. We assume that the guessing part of the program creates various combinations of ground facts for *credential*.

Each credential has a sensitivity value, all of which are fed into the external atom by the extension of *selected* for each guessed model. The weak constraint (3) causes the optimization strategy of dlhex to single out the model that has the least numerical value for *modelWeight*:

$$\begin{aligned} \text{selected}(C, V) &\leftarrow \text{credential}(C), \text{hasSens}(C, V), \\ \text{modelWeight}(X) &\leftarrow \&policy[\text{selected}](X), \\ &\Leftarrow \text{modelWeight}(X) [X : 1]. \end{aligned} \quad (3)$$

Eventually, it is up to the author of the external atom how to compute this value within the evaluation function of the *&policy*-atom. We offer a template for this plug-in, where only the actual function for the computation of the cost value needs to be inserted.

On <http://www.kr.tuwien.ac.at/research/dlhex/>, we provide a Web-interface to evaluate HEX-programs online, along with a more detailed documentation of all available external atoms. Currently, dlhex and the presented plug-ins are publicly available as source packages. Moreover, we also supply a tool kit for developing custom plug-ins, embedded in the GNU autotools environment, which takes care for the low-level, system-specific build process and lets the plug-in author concentrate his or her efforts on the implementation of the plug-in's actual core functionality.

References

1. H. Boley, S. Tabet, and G. Wagner. Design Rationale for RuleML: A Markup Language for Semantic Web Rules. In *Proc. SWWS 2001*, pages 381–401, 2001.
2. D. Calvanese, G. D. Giacomo, and M. Lenzerini. A Framework for Ontology Integration. In *Proc. SWWS 2001*, pages 303–316, 2001.
3. T. Eiter, G. Ianni, R. Schindlauer, and H. Tompits. A Uniform Integration of Higher-Order Reasoning and External Evaluations in Answer Set Programming. In *Proc. IJCAI 2005*. Morgan Kaufmann, 2005.
4. W. Faber, N. Leone, and G. Pfeifer. Recursive Aggregates in Disjunctive Logic Programs: Semantics and Complexity. In *Proc. JELIA 2004*, pages 200–212, 2004.
5. M. Gelfond and V. Lifschitz. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing*, 9:365–385, 1991.
6. V. Haarslev and R. Möller. RACER System Description. In *Proc. IJCAR 2001*, pages 701–705, 2001.
7. G. Ianni, G. Ielpa, A. Pietramala, M. C. Santoro, and F. Calimeri. Enhancing Answer Set Programming with Templates. In *Proc. NMR 2004*, pages 233–239, 2004.
8. M. Kifer, G. Lausen, and J. Wu. Logical Foundations of Object-Oriented and Frame-Based Languages. *Journal of the ACM*, 42(4):741–843, 1995.
9. N. Leone, G. Pfeifer, W. Faber, T. Eiter, G. Gottlob, S. Perri, and F. Scarcello. The DLV System for Knowledge Representation and Reasoning. *ACM Transactions on Computational Logic*. To appear.
10. M. Sintek and S. Decker. TRIPLE - A Query, Inference, and Transformation Language for the Semantic Web. In *Proc. ISWC 2002*, pages 364–378, 2002.
11. K. Wang, G. Antoniou, R. W. Topor, and A. Sattar. Merging and Aligning Ontologies in dl-Programs. In *Proc. RuleML 2005* pages 160–171, 2005.